

CLAIM AMENDMENTS

Claims 1-3. Canceled

1 4. **(Allowed)** In a computer system that includes:

2 at least one hardware processor that has a hardware instruction pointer; and

3 a first operating system (COS) initially installed to run on the hardware processor

4 at a most-privileged, system level, the system level being defined as an operational

5 state with permission to directly access predetermined physical resources of the

6 computer;

7 a method comprising:

8 initializing the computer with the COS;

9 loading a load call module within the COS;

10 upon initialization of the computer, calling a loading module from the load call

11 module, whereupon the loading module loads a kernel, which forms a second operating

12 system;

13 via the COS starting execution of the kernel, the kernel thereupon substantially

14 displacing the COS from the system level and itself running at the system level;

15 the step of loading the kernel including:

16 via the loading module, setting the hardware instruction pointer and forwarding of

17 interrupts and faults generated by the processor and by predetermined ones of the

18 system resources to point into a memory address space allocated to and controlled by

19 the kernel, the kernel thereby handling requests for system resources.

1 5. **(Allowed)** A method as in claim 4, further including the following steps:

2 after initialization of the computer, transferring from the COS to the kernel a list

3 of devices initially controlled by the COS, the devices being included among the system

4 resources; and

5 classifying the devices and control of the devices into the following groups (which

6 may be empty):

7 host-managed devices, which are controlled by the COS;

8 reserved devices, which are controlled by the kernel;

9 and shared devices, which may be controlled by either the COS or the
10 kernel.

1 6. **(Allowed)** A method as defined in claim 5, further comprising including a
2 mass storage controller as one of the shared devices.

1 7. **(Allowed)** A method as defined in claim 6, in which the mass storage
2 controller is a SCSI device.

1 8. **(Allowed)** A method as defined in claim 5, further comprising including a
2 network adapter as one of the shared devices.

1 9. **(Allowed)** A method as defined in claim 5, further comprising the steps of
2 forwarding interrupts generated by host-managed devices to the COS via the kernel,
3 and handling such interrupts within the COS.

1 10. **(Allowed)** A method as defined in claim 9, further including the step of
2 delaying handling of interrupts that are forwarded to the COS and that are generated by
3 host-managed devices until a subsequent instance of running of the COS.

1 11. **(Allowed)** A method as defined in claim 10, further including the step, upon
2 sensing, in the kernel, an interrupt raised by any host-managed device, of masking the
3 interrupt until the subsequent instance of running of the COS, thereby avoiding multiple
4 recurrences of the interrupt.

Claims 12. Canceled

1 13. **(Allowed)** In a computer system that includes:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer;
7 a method comprising:
8 initializing the computer with the COS;
9 via the COS, loading and starting execution of a kernel, which forms a second
10 operating system, the kernel thereupon substantially displacing the COS from the
11 system level and itself running at the system level, the kernel thereby handling requests
12 for system resources;
13 installing at least one virtual machine (VM) to run on the kernel via a virtual
14 machine monitor (VMM);
15 in the kernel, separately scheduling the execution of the COS and of each VM,
16 the COS and the VMs thereby forming separately schedulable and separately executing
17 entities; and
18 within the kernel, representing each schedulable entity as a corresponding world,
19 each world comprising a world memory region with a respective world address space
20 and storing a respective world control thread.

1 14. **(Allowed)** A method as defined in claim 13, further including the step of
2 switching worlds, which step comprises:
3 under control of the kernel, storing current state data for a currently executing
4 schedulable entity in a kernel-controlled memory region;
5 disabling exceptions;
6 loading state data for a subsequently executing schedulable entity;
7 starting execution of the subsequently executing schedulable entity; and
8 enabling exceptions.

1 15. **(Allowed)** A method as defined in claim 14, in which the state data for
2 each schedulable entity includes exception flags, memory segments, an instruction
3 pointer, and descriptor tables, which are loaded into an exception flags register,
4 memory segment registers, an instruction pointer register, and descriptor tables,
5 respectively.

1 16. **(Allowed)** A method as defined in claim 14, in which the computer includes
2 a plurality of hardware processors, further including the following steps:
3 in the kernel, separately scheduling the execution of each processor, the
4 processors thereby also forming separately schedulable entities;
5 within the kernel, representing each processor as a corresponding system world,
6 each system having a respective system world address space and a respective system
7 world control thread.

1 17. **(Allowed)** A method as defined in claim 16, further including the step of
2 allocating, for each processor, a separate memory mapping cache.

1 18. **(Allowed)** A method as defined in claim 13, in which each VM includes a
2 virtual processor, a virtual operating system (VOS), and an I/O driver, loaded within the
3 VOS, for an I/O device, the method further comprising the following steps:
4 allocating a shared memory space that is addressable by both the kernel and the
5 VM's I/O driver,
6 transferring an output set of data from the VM to the I/O device according to the
7 following sub-steps:
8 via the VM's I/O driver, setting a pointer to the output set of data in the
9 shared memory region and generating a request for transmission;
10 in the kernel, upon sensing the request for transmission:
11 retrieving the output set of data from a position in the shared memory region
12 indicated by the pointer and transferring the retrieved output set of data to a physical
13 transmit buffer portion of the shared memory region;
14 transferring the output data set from the physical transmit buffer portion to the
15 I/O device;

16 transferring an input set of data from the I/O device to the VM according to the
17 following sub-steps:
18 in the kernel,
19 copying the input set of data into a physical receive buffer portion
20 of the shared memory region;
21 setting the pointer to the physical receive buffer portion;
22 issuing to the VMM an instruction to raise an interrupt;
23 in the VM, upon sensing the interrupt raised by the VMM, retrieving the
24 input set of data from the physical receive buffer portion of the shared memory region
25 indicated by the pointer;
26 whereby the input and output data sets may be transferred between the VM and
27 the I/O device via only the kernel.

1 19. **(Allowed)** A method as defined in claim 18, further comprising completing
2 the sub-steps for transferring the output set of data upon sensing only a single request
3 for transmission.

1 20. **(Allowed)** A method as defined in claim 18, in which:
2 the I/O device is a network connection device for data transfer to and from a
3 network; and
4 the input and output data sets are network packets.

Claims 21-24. Canceled

1 25. **(Allowed)** In a computer system that includes:
2 at least one hardware processor, which has a hardware instruction pointer;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer;
7 a method comprising:
8 initializing the computer with the COS;
9 via the COS, loading and starting execution of a kernel, which forms a second
10 operating system, the kernel thereupon substantially displacing the COS from the
11 system level and itself running at the system level, the kernel thereby handling requests
12 for system resources;
13 halting execution of the kernel;
14 reinstating a state of the first operating system that existed before the loading of
15 the kernel; and
16 resuming execution of the first operating system at the most-privileged system
17 level;
18 the kernel thereby being functionally removed from the computer;
19 in which:
20 A) the step of loading the kernel includes the following sub-steps:
21 i) loading a load call module within the COS;
22 ii) upon initialization of the computer, calling a loading module from the
23 load call module, whereupon the loading module loads the kernel;
24 iii) after initialization of the computer, transferring from the COS to the
25 kernel a list of devices initially controlled by the COS; and
26 iv) classifying the devices and control of the devices into the following
27 groups (which may be empty):
28 a) host-managed devices, which are controlled by the COS;
29 b) reserved devices, which are controlled by the kernel; and
30 c) shared devices, which may be controlled by either the COS or
31 the kernel;

32 v) via the loading module, setting the hardware instruction pointer and
33 forwarding of interrupts and faults generated by the processor and by predetermined
34 ones of the physical resources to point into a memory address space allocated to and
35 controlled by the kernel;

36 B) the step of reinstating the state of the first operating system includes the
37 following steps:

38 i) restoring interrupt and fault handling from the kernel to the first
39 operating system;

40 ii) transferring control of host-managed and shared devices from the
41 kernel to the first operating system; and

42 iii) removing the kernel from an address space of the first operating
43 system.

1 26. **(Allowed)** A method for managing resources in a computer, which
2 includes at least one processor that has a hardware instruction pointer, the method
3 comprising the following steps:

4 A) initializing the computer using a first operating system (COS), the COS itself
5 running at a most-privileged, system level, the system level being defined as an
6 operational state with permission to directly access predetermined physical resources
7 of the computer, the physical resources including physical devices;

8 B) loading a kernel via the COS, the kernel forming a second operating system,
9 this step of loading the kernel comprising:

10 i) loading a load call module within the COS;

11 ii) upon initialization of the computer, calling a loading module from the
12 load call module, whereupon the loading module loads the kernel;

13 iii) via the loading module, setting the hardware instruction pointer and
14 forwarding interrupts and faults generated by the processor and by predetermined ones
15 of the physical resources to point into a memory address space allocated to and
16 controlled by the kernel;

17 C) starting execution of the kernel, the kernel thereupon substantially displacing
18 the COS from the system level and itself running at the system level; and

19 D) submitting requests for system resources via the kernel;

20 E) after initialization of the computer, transferring from the COS to the kernel a
21 list of the devices initially controlled by the COS; and

22 F) classifying the devices and control of the devices into the following groups
23 (which may be empty):

24 i) host-managed devices, which are controlled by the COS;
25 ii) reserved devices, which are controlled by the kernel; and
26 iii) shared devices, which may be controlled by either the COS or the
27 kernel; and

28 G) forwarding interrupts generated by host-managed devices to the COS via the
29 kernel, and handling such interrupts within the COS.

1 27. **(Allowed)** A method for managing resources in a computer, which includes
2 at least one processor that has a hardware instruction pointer, the method comprising
3 the following steps:

4 A) initializing the computer using a first operating system (COS), the COS itself
5 running at a most-privileged, system level, the system level being defined as an
6 operational state with permission to directly access predetermined physical resources
7 of the computer, the physical resources including physical devices;

8 B) loading a kernel via the COS, the kernel forming a second operating system,
9 this step of loading the kernel comprising:

10 i) loading a load call module within the COS;
11 ii) upon initialization of the computer, calling a loading module from the
12 load call module, whereupon the loading module loads the kernel;
13 iii) via the loading module, setting the hardware instruction pointer and
14 forwarding interrupts and faults generated by the processor and by predetermined ones
15 of the physical resources to point into a memory address space allocated to and
16 controlled by the kernel;

17 C) starting execution of the kernel, the kernel thereupon substantially displacing
18 the COS from the system level and itself running at the system level; and

19 D) submitting requests for system resources via the kernel;

20 E) after initialization of the computer, transferring from the COS to the kernel a
21 list of the devices initially controlled by the COS; and

22 F) classifying the devices and control of the devices into the following groups
23 (which may be empty):

- 24 i) host-managed devices, which are controlled by the COS;
- 25 ii) reserved devices, which are controlled by the kernel; and
- 26 iii) shared devices, which may be controlled by either the COS or the
27 kernel; and

28 G) forwarding interrupts generated by host-managed devices to the COS via the
29 kernel, and handling such interrupts within the COS;

30 H) installing at least one virtual machine (VM) to run on the kernel via a virtual
31 machine monitor (VMM); and

32 I) in the kernel, separately scheduling the execution of the COS and of each
33 VM, the COS and the VMs thereby forming separately schedulable and separately
34 executing entities.

Claims 28-29. Canceled

1 30. **(Allowed)** A computer system comprising:
2 at least one hardware processor that has a hardware instruction pointer;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS initializing the computer;
7 a kernel that forms a second operating system;
8 a loader comprising computer-executable code for loading the kernel via the
9 COS and for starting execution of the kernel, the kernel thereupon substantially
10 displacing the COS from the system level and itself running at the system level;
11 the kernel including a software module for handling requests for system
12 resources;
13 the loader being further provided for setting the hardware instruction pointer and
14 forwarding interrupts and faults generated by the processor and by predetermined ones
15 of the system resources to point into a memory address space allocated to and
16 controlled by the kernel.

1 31. **(Allowed)** A system as in claim 30, in which:
2 the system resources include devices initially controlled by the COS;
3 the loader is further provided for transferring, after initialization of the computer,
4 from the COS to the kernel a list of the devices initially controlled by the COS and for
5 classifying the devices and control of the devices into the following groups (which may
6 be empty):
7 host-managed devices, which are controlled by the COS;
8 reserved devices, which are controlled by the kernel; and
9 shared devices, which may be controlled by either the COS or the kernel.

1 32. **(Allowed)** A system as in claim 31, in which at least one of the shared
2 devices is a mass storage controller.

33. Canceled

1 34. **(Allowed)** A computer system comprising:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS forming means for initializing the computer;
7 a kernel that forms a second operating system;
8 a loader comprising computer-executable code for loading the kernel means via
9 the COS and for starting execution of the kernel means, the kernel thereupon
10 substantially displacing the COS from the system level and itself running at the system
11 level;
12 a virtual machine monitor (VMM);
13 at least one virtual machine (VM) installed to run on the kernel via the VMM;
14 the kernel including software modules for handling requests for system
15 resources;

16 for separately scheduling the execution of the COS and of each VM, the
17 COS and the VMs thereby forming separately schedulable and separately executing
18 entities; and
19 for representing each schedulable entity as a corresponding world, each
20 world comprising a world memory region with a respective world address space and
21 storing a respective world control thread.

1 35. **(Allowed)** A system as defined in claim 34, in which the kernel is further
2 provided:

3 for storing current state data for a currently executing schedulable entity in a
4 kernel-controlled memory region;
5 for disabling exceptions;
6 for loading state data for a subsequently executing schedulable entity;
7 for starting execution of the subsequently executing schedulable entity; and
8 for enabling exceptions;
9 the kernel thereby being provided for switching worlds.

1 36. **(Allowed)** A system as defined in claim 35, in which the state data for each
2 schedulable entity includes exception flags, memory segments, an instruction pointer,
3 and descriptor tables, which are loaded into an exception flags register, memory
4 segment registers, an instruction pointer register, and descriptor tables, respectively.

1 37. **(Allowed)** A system as defined in claim 34, in which:

2 the computer includes a plurality of hardware processors;

3 the kernel is further provided:

4 for separately scheduling the execution of each processor, the processors
5 thereby also forming separately schedulable entities;

6 for representing each processor as a corresponding system world, each
7 system having a respective system world address space and a respective system world
8 control thread.

1 38. **(Allowed)** A system as defined in claim 37, further comprising a separate
2 memory mapping cache for each processor.

1 39. **(Allowed)** A computer system comprising:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS forming means for initializing the computer;
7 a kernel that forms a second operating system;
8 a loader comprising computer-executable code for loading the kernel means via
9 the COS and for starting execution of the kernel means, the kernel thereupon
10 substantially displacing the COS from the system level and itself running at the system
11 level;
12 the kernel handling requests for system resources;
13 a virtual machine monitor (VMM);
14 at least one virtual machine (VM) installed to run on the kernel via the VMM;
15 within each VM, a virtual processor, a virtual operating system (VOS), and an I/O
16 driver for an I/O device loaded within the VOS;
17 a shared memory space that is addressable by both the kernel and the VM's I/O
18 driver, the shared memory space storing input data and output data for transfer
19 between the VM and the I/O device;
20 in which:
21 the VM's I/O driver comprises computer-executable code for setting a pointer to
22 output data in the shared memory region and generating a request for transmission;
23 the kernel is further provided, upon sensing the request for transmission:
24 for retrieving the output data from a position in the shared memory region
25 indicated by the pointer and transferring the retrieved output data to a physical transmit
26 buffer portion of the shared memory region;
27 for outputting the output data from the physical transmit buffer portion to
28 the I/O device;
29 for receiving the input data from the I/O device;

30 for copying the input data into a physical receive buffer portion of the
31 shared memory region;
32 for setting the pointer to the physical receive buffer portion;
33 for issuing to the VMM an instruction to raise an interrupt;
34 the VM's I/O driver is further provided, upon sensing the interrupt raised by the
35 VMM, for retrieving the input data from the physical receive buffer portion of the shared
36 memory region indicated by the pointer;
37 whereby the input and output data may be transferred between the VM and the
38 I/O device via only the kernel.

1 40. **(Allowed)** A system as defined in claim 39, in which:
2 the I/O device is a network connection device for data transfer to and from a
3 network; and
4 the input and output data are network packets.

41-43. Canceled

1 44. **(Allowed)** A computer system comprising:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS forming means for initializing the computer;
7 a kernel means that forms a second operating system;
8 loading means for loading the kernel via the COS and for starting execution of
9 the kernel, the kernel thereupon substantially displacing the COS from the system level
10 and itself running at the system level;
11 the kernel means is provided for handling requests for system resources;
12 in which:
13 the processor has a hardware instruction pointer;
14 the loading means is further provided for setting the hardware instruction pointer
15 and forwarding interrupts and faults generated by the processor and by predetermined

ones of the system resources to point into a memory address space allocated to and controlled by the kernel means;

the system resources include devices initially controlled by the COS;

the loading means is further provided for transferring, after initialization of the computer, from the COS to the kernel means a list of the devices initially controlled by the COS and for classifying the devices and control of the devices into the following groups (which may be empty):

- host-managed devices, which are controlled by the COS;
- reserved devices, which are controlled by the kernel means; and
- shared devices, which may be controlled by either the COS or the kernel means.

45. (Allowed) A computer system comprising:

- at least one hardware processor;
- a first operating system (COS) initially installed to run on the hardware processor at a most-privileged, system level, the system level being defined as an operational state with permission to directly access predetermined system resources of the computer, the COS forming means for initializing the computer;
- a kernel means that forms a second operating system;
- loading means for loading the kernel means via the COS and for starting execution of the kernel means, the kernel means thereupon substantially displacing the COS from the system level and itself running at the system level;
- at least one virtual machine (VM); and
- a virtual machine monitor (VMM);

in which:

- the VM is installed to run on the kernel means via the VMM;
- the kernel means is provided for handling requests for system resources;
- the processor has a hardware instruction pointer;
- the loading means is further provided for setting the hardware instruction pointer and forwarding interrupts and faults generated by the processor and by predetermined ones of the system resources to point into a memory address space allocated to and controlled by the kernel means;

21 the system resources include devices initially controlled by the COS;
22 the loading means is further provided for transferring, after initialization of the
23 computer, from the COS to the kernel means a list of the devices initially controlled by
24 the COS and for classifying the devices and control of the devices into the following
25 groups (which may be empty):
26 host-managed devices, which are controlled by the COS;
27 reserved devices, which are controlled by the kernel means; and
28 shared devices, which may be controlled by either the COS or the kernel
29 means;
30 for separately scheduling the execution of the COS and of each VM, the COS
31 and the VMs thereby forming separately schedulable and separately executing entities;
32 and
33 the kernel means is further provided:
34 for representing each schedulable entity as a corresponding world, each
35 world comprising a world memory region with a respective world address space and
36 storing a respective world control thread;
37 for storing current state data for a currently executing schedulable entity in
38 a kernel means-controlled memory region;
39 for disabling exceptions;
40 for loading state data for a subsequently executing schedulable entity;
41 for starting execution of the subsequently executing schedulable entity;
42 and
43 for enabling exceptions;
44 the kernel means thereby being provided for switching worlds.